

Script generated by TTT

Title: Petter: Programmiersprachen (06.11.2019)

Date: Wed Nov 06 12:20:06 CET 2019

Duration: 89:54 min

Pages: 26

Atomic Executions




A concurrent program consists of several threads that share *resources*:

- resources can be *memory locations* or *memory mapped I/O*
 - a file can be modified through a shared handle, e.g.
- usually *invariants* must be retained wrt. resources
 - e.g. a head and tail pointer must delimit a linked list
 - an invariant may span *multiple* resources
 - during an update, the invariant may be temporarily *locally broken*

↔ multiple resources must be updated together to ensure the invariant

Ideally, a sequence of operations that update shared resources should be *atomic* [Harris et al.(2010)Harris, Larus, and Rajwar]. This would ensure that the invariant never seems to be broken.



Definition (Atomic Execution)


A computation forms an *atomic execution* if its effect can only be *observed* as a single transformation on the memory.

Why Memory Barriers are not Enough

Often, *multiple memory locations* may only be modified exclusively by one thread during a computation.

- use barriers to implement automata that ensure *mutual exclusion*
- ↔ generalize the re-occurring *concept* of enforcing mutual exclusion

Needed: interaction with *multiple memory locations* within a *single step*:



Overview



We will address the *established* ways of managing synchronization. The presented techniques

- are available on most platforms
- likely to be found in most existing (concurrent) software
- provide solutions to common concurrency tasks
- are the source of common concurrency problems

The techniques are applicable to C, C++ (pthread), Java, C# and other imperative languages.

Wait-Free Updates

Which operations on a CPU are atomic? (j,k and tmp are registers)

Program 1

```
i++;
```

Program 2

```
j = i;
i = i+k;
```

Program 3

```
int tmp = i;
i = j;
j = tmp;
```

Answer:

- none by default (even without store and invalidate buffers, *why?*)
- ⚠ The load and store (even `i++`'s) may be interleaved with a store from another processor.

All of the programs *can* be made atomic executions (e.g. on x86):

- `i` must be in memory
- *Idea: lock the cache bus* for an address for the duration of an instruction

Program 1

```
lock inc [addr_i]
```


Program 2 (fetch-and-add)

```
mov eax,reg_k
lock xadd [addr_i],eax
mov reg_j,eax
```

Program 3 (atomic-exchange)

```
lock xchg [addr_i],reg_j
```

Wait-Free Atomic Executions



Wait-Free Bumper-Pointer Allocation

Garbage collectors often use a *bumper pointer* to allocated memory:

Bumper Pointer Allocation

```
char heap[2^20];
char* firstFree = &heap[0];

char* alloc(int size) {
    char* start = firstFree;
    firstFree = firstFree + size;

    if (start+size>sizeof(heap)) garbage_collect();
    return start;
}
```

- `firstFree` points to the first unused byte
- each allocation reserves the next `size` bytes in `heap`

Wait-Free Bumper-Pointer Allocation

Garbage collectors often use a *bumper pointer* to allocated memory:

Bumper Pointer Allocation

```
char heap[2^20];
char* firstFree = &heap[0];

char* alloc(int size) {
    char* start;
    asm("lock; xadd %0, %1" : "=r"(start), "=m"(firstFree):
        "0"(size), "m"(firstFree) : "memory");
    if (start+size>sizeof(heap)) garbage_collect();
    return start;
}
```

- `firstFree` points to the first unused byte
- each allocation reserves the next `size` bytes in `heap`

Thread-safe implementation:

- `alloc`'s core functionality matches Program 2: fetch-and-add
- ⚠ inline assembler (GCC/AT&T syntax in the example)

Marking Statements as Atomic



Rather than writing assembler: use *made-up* keyword `atomic`:

Program 1

```
atomic {
  i++;
}
```


Program 2

```
atomic {
  j = i;
  i = i+k;
}
```

Program 3

```
atomic {
  int tmp = i;
  i = j;
  j = tmp;
}
```

The statements in an `atomic` block execute as *atomic execution*:



- `atomic` only translatable when a corresponding atomic CPU instruction exist
- the notion of requesting *atomic execution* is a general concept

Lock-Free Algorithms



If a *wait-free* implementation is not possible, a *lock-free* implementation might still be viable.

Common usage pattern for *compare and swap*:

- 1 read the initial value in i into k (using memory barriers)
- 2 compute a new value $j = f(k)$
- 3 update i to j if $i = k$ still holds
- 4 go to first step if $i \neq k$ meanwhile

$$i = f(i)$$

Wait-Free Synchronization



Wait-Free algorithms are limited to a single instruction:

- no control flow possible, no behavioral change depending on data
- often, there are instructions that execute an operation conditionally

Program 4

```
atomic {
  r = b;
  b = 0;
}
```

Program 5

```
atomic {
  r = b;
  b = 1;
}
```

Program 6

```
atomic {
  r = (k==i);
  if (r) i = j;
}
```

Operations *update* a memory cell and *return* the previous value.

- the first two operations can be seen as setting a flag b to $v \in \{0, 1\}$ and returning its previous state.
 - ▶ the operation implementing programs 4 and 5 is called *set-and-test*
- the third case generalizes this to setting a variable i to the value of j , if i 's old value is equal to k 's.
 - ▶ the operation implementing program 6 is called *compare-and-swap*

↪ use as *building blocks* for algorithms that can *fail*

Lock-Free Algorithms



If a *wait-free* implementation is not possible, a *lock-free* implementation might still be viable.

Common usage pattern for *compare and swap*:

- 1 read the initial value in i into k (using memory barriers)
- 2 compute a new value $j = f(k)$
- 3 update i to j if $i = k$ still holds
- 4 go to first step if $i \neq k$ meanwhile

⚠ note: $i = k$ must imply that no thread has updated i

General recipe for *lock-free* algorithms

- given a *compare-and-swap* operation for n bytes
- try to *group variables for which an invariant must hold* into n bytes
- *read these bytes atomically*
- compute a new value
- perform a *compare-and-swap* operation on these n bytes

Locked Atomic Executions

Semaphores and Mutexes

A (counting) *semaphore* is an integer s with the following operations:



```
void signal(int *s) {
    atomic { *s = *s + 1; }
}
```

```
void wait(int *s) {
    bool avail;
    do {
        atomic {
            avail = *s>0;
            if (avail) (*s)--;
        }
    } while (!avail);
}
```

A counting semaphore can track how many resources are still available.

- a thread *acquiring* a resource executes `wait()`
- if a resource is still available, `wait()` returns
- once a thread finishes using a resource, it calls `signal()` to *release*

Special case: initializing with `s = 1` gives a *binary* semaphore:

- can be used to block and unblock a thread
- can be used to protect a single resource

↔ in this case the data structure is also called `mutex`

Locks



Definition (Lock)

A lock is a data structure that

- can be *acquired* and *released*
- ensures *mutual exclusion*: only one thread may hold the lock at a time
- *blocks* other threads attempts to acquire while held by a different thread
- protects a *critical section*: a piece of code that may produce incorrect results when entered concurrently from several threads

⚠ may *deadlock* the program

Implementation of Semaphores

A *semaphore* does not have to wait busily:



```
void signal(int *s) {
    atomic { *s = *s + 1; }
    wake(s);
}
```

```
void wait(int *s) {
    bool avail;
    do {
        atomic {
            avail = *s>0;
            if (avail) (*s)--;
        }
    } while (!avail);
    de_schedule(s);
}
```

Busy waiting is avoided:

- a thread failing to decrease $*s$ executes `de_schedule()`
- `de_schedule()` enters the operating system and inserts the current thread into a queue of threads that will be woken up when $*s$ becomes non-zero, usually by *monitoring writes* to s (↔ *FUTEX_WAIT*)
- once a thread calls `wake(s)`, the first thread t waiting on s is extracted
- the operating system lets t return from its call to `de_schedule()`

Practical Implementation of Semaphores



Certain optimisations are possible:



```
void signal(int *s) {
    atomic { *s = *s + 1; }
    wake(s);
}
```

```
void wait(int *s) {
    bool avail;
    do {
        atomic {
            avail = *s>0;
            if (avail) (*s)--;
        }
        if (!avail) de_schedule(s);
    } while (!avail);
}
```

In general, the implementation is more complicated

- `wait()` may busy wait for a few iterations
 - ▶ avoids de-scheduling if the lock is released frequently
 - ▶ better throughput for semaphores that are held for a short time
- `wake(s)` informs the scheduler that `s` has been written to

~ using a semaphore with a single core reduces to

```
if (*s) (*s)--; /* critical section */ (*s)++;
```

Implementation of a Basic Monitor



A monitor contains a semaphore `count` and the id `tid` of the occupying thread:

```
typedef struct monitor mon_t;
struct monitor { int tid; int count; };
void monitor_init(mon_t* m) { memset(m, 0, sizeof(mon_t)); }
```

Define `monitor_enter` and `monitor_leave`:

- ensure mutual exclusion of accesses to `mon_t`
- track how many times we called a monitored procedure recursively

```
void monitor_enter(mon_t *m) {
    bool mine = false;
    while (!mine) {
        mine = thread_id()==m->tid;
        if (mine) m->count++; else
            atomic {
                if (m->tid==0) {
                    m->tid = thread_id();
                    mine = true; m->count=1;
                }
            };
        if (!mine) de_schedule(&m->tid);
    }
}

void monitor_leave(mon_t *m) {
    m->count--;
    if (m->count==0) {
        atomic {
            m->tid=0;
        }
        wake(&m->tid);
    }
}
```

Monitors: An Automatic, Re-entrant Mutex



Often, a data structure can be made thread-safe by

- *acquiring* a lock upon *entering* a function of the data structure
- *releasing* the lock upon *exit* from this function

Locking each procedure body that accesses a data structure:

- 1 is a re-occurring pattern, should be generalized
- 2 becomes problematic in recursive calls: it blocks

E.g. a thread `t` waits for a data structure to be filled

- ▶ `t` will call `pop()` and obtain `-1`
- ▶ `t` then has to call again, until an element is available

~ `t` is busy waiting and produces contention on the lock ⚠



Monitor: a mechanism to address these problems:

- 1 a procedure associated with a monitor **acquires a lock** on entry and **releases** it on exit
- 2 if that lock is already taken by the current thread, proceed

Condition Variables



✓ Monitors simplify the construction of thread-safe resources.

Still: Efficiency problem when using resource to synchronize:

E.g. a thread `t` waits for a data structure to be filled:

- ▶ `t` will call `pop()` and obtain `-1`
- ▶ `t` then has to call again, until an element is available

~ `t` is busy waiting and produces contention on the lock

Idea: create a *condition variable* on which to block while waiting:

```
struct monitor { int tid; int count; int cond; int cond2;... };
```

Condition Variables

✓ Monitors simplify the construction of thread-safe resources.

Still: Efficiency problem when using resource to synchronize:

E.g. a thread t waits for a data structure to be filled:

- ▶ t will call `pop()` and obtain `-1`
- ▶ t then has to call again, until an element is available
- ~> t is busy waiting and produces contention on the lock

Idea: create a *condition variable* on which to block while waiting:


```
struct monitor { int tid; int count; int cond; int cond2;... };
```

Define these two functions:

- 1 **wait** for the condition to become true
 - ▶ called while being *inside* the monitor
 - ▶ temporarily *releases* the monitor and blocks
 - ▶ when *signalled*, re-acquires the monitor and returns
- 2 **signal** waiting threads that they may be able to proceed
 - ▶ one/all waiting threads that called *wait* will be woken up, two possibilities:
 - signal-and-urgent-wait** : the *signalling* thread suspends and continues once the *signalled* thread has released the monitor
 - signal-and-continue** the *signalling* thread continues, any *signalled* thread enters when the monitor becomes available

Signal-And-Continue Semantics

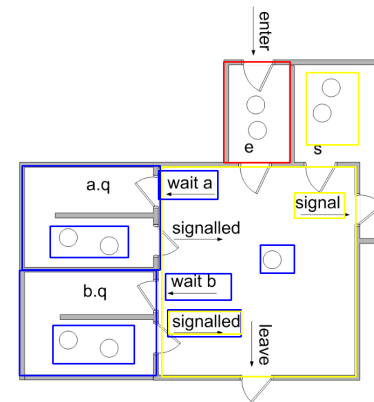
Here, the `signal` function is usually called `notify`.



- a call to `wait` on condition a adds thread to the queue $a.q$
- a call to `notify` for a adds one thread from $a.q$ to e (unless $a.q$ is empty)
- if a thread leaves, it wakes up one thread waiting on e
- ~> signalled threads compete for the monitor
- assuming FIFO ordering on e , threads who tried to enter between `wait` and `notify` will run first
- need additional queue s if waiting threads should have priority

Signal-And-Urgent-Wait Semantics

Requires one queue for each condition c and a suspended queue s :



source: [http://en.wikipedia.org/wiki/Monitor_\(synchronization\)](http://en.wikipedia.org/wiki/Monitor_(synchronization))

- a thread who tries to enter a monitor is added to queue e if the monitor is occupied
- a call to `wait` on condition a adds thread to the queue $a.q$
- a call to `signal` for a adds thread to queue s (suspended)
- one thread from the a queue is woken up
- `signal` on a is a no-op if $a.q$ is empty
- if a thread leaves, it wakes up one thread waiting on s
- if s is empty, it wakes up one thread from e

~> queue s has priority over e

Implementing Condition Variables

We implement the simpler *signal-and-continue* semantics for a single condition variable:

~> a *notified* thread is simply woken up and competes for the monitor

```
void cond_wait(mon_t *m) {
    assert(m->tid==thread_id());
    int old_count = m->count;
    m->tid = 0;
    wait(&m->cond);
    bool next_to_enter;
    do {
        atomic {
            next_to_enter = m->tid==0;
            if (next_to_enter) {
                m->tid = thread_id();
                m->count = old_count;
            }
        }
        if (!next_to_enter) de_schedule(&m->tid);
    } while (!next_to_enter);}
}
```

```
void cond_notify(mon_t *m) {
    // wake up other threads
    signal(&m->cond);
}
```

A Note on Notify

With *signal-and-continue* semantics, two notify functions exist:

- 1 `notify`: wakes up exactly one thread waiting on condition variable
- 2 `notifyAll`: wakes up all threads waiting on a condition variable

⚠ an implementation often becomes easier if `notify` means *notify some*


~ programmer should assume that thread is not the only one woken up

Deadlocks



Monitors with a Single Condition Variable

Monitors with a single condition variable are built into *Java* and *C#*:



source: [http://en.wikipedia.org/wiki/Monitor_\(synchronization\)](http://en.wikipedia.org/wiki/Monitor_(synchronization))

```
class C {  
    public synchronized void f() {  
        // body of f  
    }  
}
```

is equivalent to

```
class C {  
    public void f() {  
        monitor_enter(this);  
        // body of f  
        monitor_leave(this);  
    }  
}
```

with `Object` containing:

```
private int mon_var;  
private int mon_count;  
private int cond_var;  
protected void monitor_enter();  
protected void monitor_leave();
```

